

Distinguishing and integrating aleatoric and epistemic variation in uncertainty quantification

Kamaljit Chowdhary* and Paul Dupuis†

January 14, 2013

Abstract

Much of uncertainty quantification to date has focused on determining the effect of variables modeled probabilistically, and with a known distribution, on some physical or engineering system. We develop methods to obtain information on the system when the distributions of some variables are known exactly, others are known only approximately, and perhaps others are not modeled as random variables at all. The main tool used is the duality between risk-sensitive integrals and relative entropy, and we obtain explicit bounds on standard performance measures (variances, exceedance probabilities) over families of distributions whose distance from a nominal distribution is measured by relative entropy. The evaluation of the risk-sensitive expectations is based on polynomial chaos expansions, which help keep the computational aspects tractable.

1 Introduction

Uncertainty quantification refers to a broad set of techniques for understanding the impact of uncertainties in complicated mechanical and physical systems. In this context “uncertainty” can take on many meanings, but we follow the convention of dividing uncertainty into aleatoric and epistemic categories. In short, aleatoric uncertainty refers to inherent uncertainty due to stochastic or probabilistic variability. This type of uncertainty is *irreducible* in that there will always be positive variance since the underlying variables are truly random. Epistemic uncertainty refers to limited knowledge we may have about the model or system. This type of uncertainty is *reducible* in that if we have more information, e.g., take more measurements, then this type of uncertainty can be reduced. However, for many problems where uncertainty quantification

*Division of Applied Mathematics, Brown University, Providence, RI, 02912, USA. Research supported in part by the Air Force Office of Scientific Research (FA9550-07-1-0425).

†Lefschetz Center for Dynamical Systems, Division of Applied Mathematics, Brown University, Providence, RI, 02912, USA. Research supported in part by the National Science Foundation (DMS-1008331), the Department of Energy (DE-SC0002413), and the Air Force Office of Scientific Research (FA9550-07-1-0425 and FA9550-09-1-0378).

is important, the acquisition of data is difficult or expensive. The epistemic uncertainty cannot be removed entirely, and so one needs modeling and computational techniques which can also accommodate this form of uncertainty.

Much of the work to date has focused on aleatoric uncertainty and what one might call the propagation of uncertainty. Here, one characterizes uncertainty concerning one or more elements of a physical system via some probability distribution, and attempts to quantify how that uncertainty will be propagated throughout the system under its constitutive equations and laws. The simplest and most straightforward method to characterize this output distribution would be to sample from the input probability distribution, solve the system equations, and thereby produce output samples (i.e., standard Monte Carlo). Although this is very simple conceptually, it can be far from practical, especially when it is computationally intensive to construct the mapping that takes input variables to output. Straightforward schemes such as Monte Carlo would require that the system be solved for each sample of the random variable, a situation that is often not practical. In recent years efficient computational methods have been developed to calculate particular functionals of the induced distribution that may be required for a particular application (e.g., covariances or error probabilities). Most recently, polynomial chaos has become popular as an efficient method for approximating the distribution of the output variable.

We note that random variables with a well-defined and given distribution are often used in this context even when there is no justification for their use, as in the case of modeling error. In other words, it is (perhaps implicitly) assumed that epistemic uncertainty can be modeled by aleatoric uncertainty. One reason is that, at least as they have been developed to date, most computational techniques (e.g., polynomial chaos and Monte Carlo) are based on the assumption that the user can identify some “appropriate” distribution for each uncertain aspect of the system, regardless of the type of uncertainty, aleatoric or epistemic. If one is interested in just basic qualitative properties of the system then this may not be a central issue, since virtually any model of uncertainty will give information on the sensitivities of the system. However, when the intended use of uncertainty quantification is for regulatory assessment or some other application where performance measures are sensitive to distributional assumptions, the issue becomes much more important, and one should carefully distinguish how one accounts for the two types of uncertainty.

The aim of the present paper is to describe an approach that (i) logically distinguishes those aspects of uncertainty that are treated as stochastic variability from other forms of uncertainty, (ii) in cases where a stochastic model is theoretically valid but for which determination of the distribution is not practical, gives bounds for performance measures that are valid for explicitly identified families of distributions, and (iii) is computationally feasible if ordinary uncertainty propagation is feasible. The intended audience is broad, including numerical analysts interested in robust performance bounds as well as applied probabilists for whom certain computational aspects of uncertainty quantification may be novel. Since a typical reader might not be familiar with the terminology and methods from both fields, we have included background material for both the

probabilistic and numerical analysis approaches used to make the paper broadly accessible.

The paper is organized as follows. In Section 2 we further discuss the distinction between aleatoric and epistemic uncertainty, and explain some of the limitations to achieving useful performance bounds in the presence of epistemic uncertainty. Section 3 introduces risk-sensitive performance measures and discusses their robust properties. Several hybrid forms are introduced that are more useful than the simplest form when both aleatoric and epistemic uncertainties are present. Various properties of the risk-sensitive measures that are useful in applications (monotonicity, optimization) are also discussed. Section 4 presents numerical examples, reviews the computational methods used to evaluate the risk-sensitive performance measures, and finishes with a subsection of discussion and conclusions.

2 Aleatoric and epistemic uncertainties

As noted in the Introduction, *uncertainty* can be divided into two categories, aleatoric and epistemic. From the perspective of mathematical formulation and modeling, aleatoric uncertainty is in some sense simpler. In contrast, epistemic uncertainty can mean different things in different contexts. *Lack of knowledge* is an ambiguous term that encompasses many different scenarios.

To illustrate, we consider an elementary example. Consider a system described by some well-posed partial differential equation (PDE), but with an uncertain boundary condition. The boundary condition is expressed in terms of a random variable X (e.g., $u_x(0, t) = X$) with a particular fixed distribution, and the numerical analysis problem is to characterize the induced distribution of the solution to the PDE at some time and location (e.g., $u(x_0, t_0; X)$). Suppose we know that the boundary condition is properly modeled as a random variable, but we do not know the correct value of some parameter in that distribution. For example, based on a central limit type argument, one might claim the distribution is known to be Gaussian, but still unknown are the “true” mean and/or variance. In this case, the *lack of knowledge* is this missing information about the parameters of the distribution. The lack of information regarding these parameters is a form of epistemic uncertainty. Hence in this example aleatoric and epistemic uncertainty are mingled. Assuming that samples are available, one could use the empirical mean and variance from data to approximate the true mean and variance, and thus reduce this uncertainty. However, in the common situation where sampling is necessarily limited, some epistemic uncertainty is inherent in the model due to practical limitations on data acquisition and modeling.

Another type of epistemic uncertainty is the omission of important aspects of the system model. This form is possibly the most difficult to quantify. For example, there might be a hidden random variable in the model. In the PDE example we have assumed the boundary condition is modeled via a random variable, but there may be other coefficients in the model that are random but

treated as constants. It is also frequently true that the mathematical model used for the system is only approximately true. For example, the model might simplify the geometry of the true system, nonlinearities may be approximated by linear relations, etc. In the context of the PDE example, boundary conditions of Dirichlet form were chosen, when in fact a more realistic model might use a mixed form (e.g., Robin boundary conditions).

For almost all forms of epistemic uncertainty there is little justification for the use of random variables to model the uncertainty. Nonetheless, it is common practice to do just that, and in practice randomness and random variables are used in many situations to account for errors in the physical model or modeling ignorance. The reasons were mentioned previously—that basic distributional properties of the output (e.g., variance) will still provide a reasonable sensitivity analysis for the problem, and existing computational methods are largely based on aleatoric uncertainty. However, with little justification for the use of any particular distribution (or even the use of randomness at all), in more critical applications one may insist on rigorous bounds on performance that are valid for a particular family of distributions. In the extreme case where no probabilistic model is considered acceptable, one may wish to only prescribe bounds on certain parameters, and then obtain tight bounds on performance over all values of the parameters that satisfy the bounds.

Thus there are many different types of uncertainty that one should account for in an analysis of the effects of uncertainty. These include: (i) aleatoric with known distribution; (ii) aleatoric with partly known distribution (mingled aleatoric and epistemic); (iii) epistemic for which one is willing to model by a family of aleatoric uncertainties, and (iv) epistemic where one is only willing to place bounds on the uncertainties. As remarked in the Introduction, this paper will introduce an approach that allows these uncertainties to be handled within a single framework that can exploit computational methods originally developed just for the treatment of aleatoric uncertainties.

3 Duality for exponential integrals

Our development of performance measures that distinguish forms of uncertainty, and which in particular allow for robustness with respect to epistemic uncertainties, depends on a duality relation between exponential integrals and relative entropy. Its first use along these lines but with regard to estimation appears to be in [1], and for optimization in [3]. We first state the basic duality result, and then define two functionals which will allow aleatoric and epistemic uncertainties to be analyzed simultaneously but at the same time differentiated.

The general duality is stated for random variables that take values in a Polish space (i.e., a complete, separable metric space) \mathcal{X} . The associated σ -algebra is the Borel σ -algebra. A typical example of \mathcal{X} for our purposes is a closed subset of some Euclidean space \mathbb{R}^d . Let $\mathcal{P}(\mathcal{X})$ denote the collection of probability measures on \mathcal{X} , and let $\mu \in \mathcal{P}(\mathcal{X})$. Given any $\nu \in \mathcal{P}(\mathcal{X})$ that is absolutely continuous with respect to μ , we define the relative entropy of ν with respect

to μ by

$$R(\nu \parallel \mu) = \int_{\mathcal{X}} \log \left(\frac{d\nu}{d\mu}(x) \right) \nu(dx)$$

whenever $\log(d\nu/d\mu(x))$ is integrable with respect to ν . In all other cases $R(\nu \parallel \mu)$ is defined to be ∞ .

Relative entropy defines a mapping $(\nu, \mu) \rightarrow R(\nu \parallel \mu)$ from $\mathcal{P}(\mathcal{X})^2$ to $\mathbb{R} \cup \infty$. This mapping has a number of very attractive properties. For example, relative entropy is non-negative, with $R(\nu \parallel \mu) = 0$ if and only if $\nu = \mu$. In addition, the mapping is jointly convex and lower semicontinuous [2, Lemma 1.4.3]. A property of particular interest for our purposes is the following variational formula for exponential integrals, which can be considered as an infinite dimensional version of the Legendre transform. Given any bounded and continuous function $F : \mathcal{X} \rightarrow \mathbb{R}$ and any $c \in (0, \infty)$,

$$\Lambda_c \doteq \frac{1}{c} \log \int_{\mathcal{X}} e^{cF(x)} \mu(dx) = \sup_{\nu \in \mathcal{P}(\mathcal{X})} \left[-\frac{1}{c} R(\nu \parallel \mu) + \int_{\mathcal{X}} F(x) \nu(dx) \right]. \quad (1)$$

For a proof see [2, Proposition 1.4.2]. The duality continues to hold as stated if F is bounded from above, and also when bounded from below if one restricts the supremum to $\nu \in \mathcal{P}(\mathcal{X})$ for which $R(\nu \parallel \mu) < \infty$ [2, Proposition 4.5.1].

As an elementary example on how such a formula can be useful, suppose that F is a performance measure (e.g., variance or an error probability), and that μ is a model for some random phenomena. We consider μ to be the *nominal model*, e.g., our best guess. However, we are uncertain if this is the correct model, and would like a measure of performance that also holds for alternative models, but with a penalty for deviation from the nominal model. Λ_c is just such a performance measure, since by (1), for any alternative model ν the bound

$$\int_{\mathcal{X}} F(x) \nu(dx) \leq \Lambda_c + \frac{1}{c} R(\nu \parallel \mu)$$

applies. Hence we obtain bounds on performance over a family of alternative models. The parameter c allows one to balance robustness with respect to possible model inaccuracies against tighter bounds. In particular, as c tends to 0, Λ_c converges to $\int_{\mathcal{X}} F(x) \mu(dx)$, which is the performance measure under the nominal model, but in the limit the bound is meaningful only when $\nu = \mu$.

Suppose that a bound on performance over a specific family of distributions is needed. Let R^* denote the maximum of relative entropy with respect to the nominal model over this family. Then the tightest possible bound is obtained by minimizing

$$\Lambda_c + \frac{1}{c} R^*$$

over $c > 0$. We show in Proposition 3 below that this function has only one local minimum over $c \in (0, \infty]$, and thus the global minimum is easy to compute.

Functionals of the exponential form that appears in the definition of Λ_c are sometimes called *risk-sensitive* because the exponential amplifies the effect of

any large values of the performance measure F . In the next two sections we consider two different hybrid risk-sensitive functionals that will allow aleatoric and epistemic variables to be combined and yet differentiated.

3.1 Two hybrid forms

In this section and the next random variables with a *known* distribution will take values in a Polish space \mathcal{X} . Variables whose distribution is *not known* or are otherwise of the epistemic variety take values in the space \mathcal{Y} .

The performance measure of interest for some given problem is assumed to be of the form

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} F(x, y) \gamma(dy) \mu(dx), \quad (2)$$

where μ (resp., γ) is a probability measure on \mathcal{X} (resp., \mathcal{Y}). If X and Y are independent random variables with distributions μ and γ , then $F(X, Y)$ represents both the performance measure (e.g., a second moment) as well as the underlying physical or mechanical system that maps these aleatoric and epistemic inputs into outputs. The analogous ordinary risk-sensitive performance measure is

$$\Lambda_c = \frac{1}{c} \log \int_{\mathcal{X}} \int_{\mathcal{Y}} e^{cF(x, y)} \gamma(dy) \mu(dx). \quad (3)$$

Neither of the measures (2) or (3) differentiate the variables according to type (aleatoric or epistemic) and given that the performance measure of interest is actually F , the use of a risk-sensitive version of the cost is not well motivated for the aleatoric variables. Indeed, use of this measure will give bounds that are robust with respect to variations on a distribution that is known, and obviously such bounds will not be as tight as possible.

The first form we consider for a hybrid measure is

$$\Lambda_c^1 = \frac{1}{c} \log \int_{\mathcal{Y}} e^{\int_{\mathcal{X}} cF(x, y) \mu(dx)} \gamma(dy). \quad (4)$$

Note that by Jensen's inequality (applied to the convex function \exp), this is smaller than (and in general strictly smaller than) Λ_c . Using the relative entropy representation for exponential integrals given in (1) [but with \mathcal{Y} in place of \mathcal{X} and letting $\int_{\mathcal{X}} cF(x, y) \mu(dx)$ be the cost function], it follows that for any distribution $\theta(dy)$

$$\int_{\mathcal{Y}} \int_{\mathcal{X}} F(x, y) \mu(dx) \theta(dy) \leq \frac{1}{c} R(\theta(dy) \parallel \gamma(dy)) + \Lambda_c^1. \quad (5)$$

This gives a bound on the performance measure for an arbitrary distribution on Y , but with the distribution on X equal to the known true distribution. The distributions thus play very different roles. In particular, we think of γ as a *nominal* distribution of Y , which should be distinguished from a possible *true* distribution. The risk sensitive functional Λ_c^1 , whose numerical evaluation can be carried out by a variety of methods (including polynomial chaos expansion

as discussed below), is based on the nominal distribution. Through the relative entropy duality, it will yield various bounds (depending on c) on a families of distributions, with the relative entropy distance the key metric. Ideally one would consider the smallest family that contains the “true” distribution (if it exists), and get the tightest bound by optimizing over c . With the hybrid risk-sensitive formulation, μ is allowed to be both the nominal (computational) distribution and also the true distribution.

The second form one could consider for a hybrid measure is

$$\Lambda_c^2 = \frac{1}{c} \int_{\mathcal{X}} \left[\log \int_{\mathcal{Y}} e^{cF(x,y)} \gamma(dy) \right] \mu(dx). \quad (6)$$

Note that by Jensen’s inequality (applied now to the concave function \log), this is again in general strictly smaller than Λ_c . The relative entropy representation for exponential integrals gives

$$\int_{\mathcal{Y}} F(x, y) \theta(dy|x) \leq \frac{1}{c} R(\theta(dy|x) \parallel \gamma(dy)) + \frac{1}{c} \log \int_{\mathcal{Y}} e^{cF(x,y)} \gamma(dy).$$

Here $\theta(dy|x)$ is any stochastic kernel on \mathcal{Y} given \mathcal{X} , which is essentially a conditional distribution on \mathcal{Y} given $X = x$ (for the precise definition see [2, page 35]). Integrating over \mathcal{X} gives

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} F(x, y) \theta(dy|x) \mu(dx) \leq \frac{1}{c} \int_{\mathcal{X}} R(\theta(dy|x) \parallel \gamma(dy)) \mu(dx) + \Lambda_c^2.$$

In the case where $\theta(dy|x)$ is independent of x , we obtain

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} F(x, y) \theta(dy) \mu(dx) \leq \frac{1}{c} R(\theta(dy) \parallel \gamma(dy)) + \Lambda_c^2.$$

Note that the second hybrid cost gives a more general bound, in that it allows the alternative distribution on Y (i.e., $\theta(dy|x)$) to depend on the value taken by X . This additional flexibility comes at a cost, and indeed we will show in the next subsection that the following inequalities hold (typically in a strict fashion):

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} F(x, y) \gamma(dy) \mu(dx) \leq \Lambda_c^1 \leq \Lambda_c^2 \leq \Lambda_c. \quad (7)$$

Hence if one is concerned with performance bounds for distributions on Y that do not depend on the variable X , then the tightest bound is obtained using Λ_c^1 .

Before studying properties of the risk-sensitive measures, we note some elementary generalizations that are possible. Suppose that the nominal distribution of X and Y is not of product form. In the setting of the first form, we could let γ be the marginal distribution of Y and let $\mu(dx|y)$ be the conditional distribution of X given $Y = y$. Using the extended definition

$$\bar{\Lambda}_c^1 = \frac{1}{c} \log \int_{\mathcal{Y}} e^{\int_{\mathcal{X}} cF(x,y) \mu(dx|y)} \gamma(dy),$$

we obtain the bound

$$\int_{\mathcal{Y}} \int_{\mathcal{X}} F(x, y) \mu(dx|y) \theta(dy) \leq \frac{1}{c} R(\theta(dy) \parallel \gamma(dy)) + \bar{\Lambda}_c^1.$$

For the second form, we let μ be the marginal distribution of X and let $\gamma(dy|x)$ be the conditional distribution of Y given $X = x$. With the definition

$$\bar{\Lambda}_c^2 = \frac{1}{c} \int_{\mathcal{X}} \left[\log \int_{\mathcal{Y}} e^{cF(x,y)} \gamma(dy|x) \right] \mu(dx),$$

we obtain

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} F(x, y) \theta(dy|x) \mu(dx) \leq \frac{1}{c} \int_{\mathcal{X}} R(\theta(dy|x) \parallel \gamma(dy|x)) \mu(dx) + \bar{\Lambda}_c^2.$$

It is indeed sometimes useful to allow the distribution of one type of variable to depend on the value taken by the other type of variable. As an elementary example, consider the situation where an aleatoric variable appears in a system, and while it is known that this variable has a Gaussian distribution with mean zero, the variance of the variable is however not known. Then the variance itself might be modeled as an uncertainty whose distribution is not known precisely, i.e., as an epistemic uncertainty. In this case $\bar{\Lambda}_c^1$ would be the relevant risk-sensitive formulation. An example of this sort is presented in Subsection 4.2.2.

Remark 1 In problems of regulatory assessment epistemic uncertainty can pose a unique set of challenges. The collection of data and model validation are often particularly difficult or expensive, while at the same time stringent bounds on performance might be demanded. In these circumstances, the framework presented here, wherein tight bounds are obtained for a known family of alternative models, would seem very natural. Prior to the collection of data or testing of samples to determine compliance, all critical elements of the model, including selection of the nominal model, size of the family of alternative models allowed by the relative entropy bounds, and the performance measures that would be required to hold uniformly for this family, would all be determined by negotiation among the interested parties. This would allow various competing needs (e.g., expense of model validation versus adequate protection of consumers) to be balanced according to the interests of the participants.

3.2 Properties of the risk-sensitive forms

In this section we prove properties of the two hybrid risk-sensitive forms. In particular we derive an inequality relating the two, and study the limit as $c \rightarrow \infty$. For the results to hold as stated we often need F be bounded from below. This is a mild assumption. Indeed, standard measures of performance such as second moments and error probabilities satisfy this condition.

The first proposition shows the inequality between the two hybrid forms.

Proposition 2 *Assume that F is bounded from below and consider the functionals defined in (2), (3), (4) and (6). Then the inequalities (7) hold.*

Proof. The only inequality which does not follow from Jensen's inequality is $\Lambda_c^1 \leq \Lambda_c^2$. The relative entropy duality as stated in (1) applied to Λ_c^1 gives

$$\Lambda_c^1 = \sup_{\theta \in \mathcal{P}(\mathcal{Y})} \left[\int_{\mathcal{Y}} \int_{\mathcal{X}} F(x, y) \mu(dx) \theta(dy) - \frac{1}{c} R(\theta(dy) \parallel \gamma(dy)) \right],$$

where $\mathcal{P}(\mathcal{Y})$ is the space of probability measures on \mathcal{Y} . The same representation applied to

$$\frac{1}{c} \log \int_{\mathcal{Y}} e^{cF(x, y)} \gamma(dy)$$

gives

$$\frac{1}{c} \log \int_{\mathcal{Y}} e^{cF(x, y)} \gamma(dy) = \sup_{\theta \in \mathcal{P}(\mathcal{Y})} \left[\int_{\mathcal{Y}} F(x, y) \theta(dy) - \frac{1}{c} R(\theta(dy) \parallel \gamma(dy)) \right].$$

The supremum operation in the last display can be done in such a way that an optimizing (or near optimizing) θ is a Borel measurable function of x , i.e., a stochastic kernel $\theta(dy|x)$ on \mathcal{Y} given \mathcal{X} [2]. Let $\mathcal{P}(\mathcal{Y}|\mathcal{X})$ denote the space of such stochastic kernels. Integrating the last display with respect to μ gives

$$\Lambda_c^2 = \sup_{\theta \in \mathcal{P}(\mathcal{Y}|\mathcal{X})} \left[\int_{\mathcal{X}} \int_{\mathcal{Y}} F(x, y) \theta(dy|x) \mu(dx) - \frac{1}{c} \int_{\mathcal{X}} R(\theta(dy|x) \parallel \gamma(dy)) \mu(dx) \right].$$

Since $\mathcal{P}(\mathcal{Y}) \subset \mathcal{P}(\mathcal{Y}|\mathcal{X})$, $\Lambda_c^2 \geq \Lambda_c^1$ follows. ■

The next proposition shows how to obtain tight bounds on the performance over a family of distributions defined in terms of a maximum relative entropy B .

Proposition 3 *Consider the functionals defined in (3), (4) and (6), and let $D = \{c : \Lambda_c < \infty\}$ (resp., $D^i = \{c : \Lambda_c^i < \infty\}, i = 1, 2$). Assume that the interior of D (resp., D^i) is nonempty. Then Λ_c (resp., Λ_c^i) is differentiable on the interior of D (resp., D^i). Assume also that F is bounded from below by zero. Then $c \rightarrow \Lambda_c$ (resp., $c \rightarrow \Lambda_c^i$) is nondecreasing for $c \geq 0$. Let $B > 0$ be given. Then there is a unique $c \in (0, \infty]$ at which*

$$c \rightarrow \frac{1}{c} B + \Lambda_c \quad \left(\text{resp., } \frac{1}{c} B + \Lambda_c^i \right)$$

attains a local minimum, where the statement that the minimum occurs at $c = \infty$ means that $\Lambda_c + B/c > \Lambda_\infty$ for a well defined limit Λ_∞ and all $c < \infty$.

Proof. To simplify we give the proof only for the case of Λ_c and omit the y variable. Thus we consider $\Lambda_c = \frac{1}{c} \log \int_{\mathcal{X}} e^{cF(x)} \mu(dx)$. Proofs for the other functionals are analogous. It is well known that $H(c) = \log \int_{\mathcal{X}} e^{cF(x)} \mu(dx)$ is

a convex function taking values in $\mathbb{R} \cup \{\infty\}$, and strictly convex and infinitely differentiable on the interior of $\{c : H(c) < \infty\}$ (see, e.g., [4]).

Next we prove the monotonicity. Differentiating gives

$$H'(c) = \frac{\int_{\mathcal{X}} F(x) e^{cF(x)} \mu(dx)}{\int_{\mathcal{X}} e^{cF(x)} \mu(dx)}.$$

Since $F \geq 0$ the derivative is non-negative, and convexity implies it is non-decreasing. Using

$$\frac{1}{c}H(c) = \frac{1}{c}(H(c) - H(0)) = \frac{1}{c} \int_0^c H'(s) ds,$$

it follows that $\Lambda_c = H(c)/c$ is nondecreasing for $c \geq 0$.

First assume that $M \doteq \sup D = \infty$, and observe that

$$\frac{d}{dc} \left[\frac{1}{c}B + \frac{1}{c}H(c) \right] = \frac{1}{c^2} \left[-B + cH'(c) - H(c) \right].$$

Since H is strictly convex and $H'(0) \geq 0$, the mapping $f(c) \rightarrow cH'(c) - H(c)$ is monotone increasing and takes the value 0 at $c = 0$ (see Figure 1).




Figure 1: $f(c)$ is monotone increasing

Let K denote the limit of $f(c)$ as $c \rightarrow \infty$. If $K = \infty$ then there is a unique solution to

$$cH'(c) - H(c) = B,$$

and we are done. The same is true if $0 \leq B < K$. Hence the only case left is when $B \geq K$. In this case $(B+H(c))/c$ is monotone decreasing for all $c \in [0, \infty)$.

Since $H(c)/c \geq H'(0) \geq 0$, there is a well-defined limit Λ_∞ that is necessarily the minimum.

Next assume that $M \doteq \sup D \in (0, \infty)$. We claim that in this case $cH'(c) - H(c) \uparrow \infty$ as $c \uparrow M$, and therefore we can argue just as before. By monotone convergence it must be that $H(c) \uparrow \infty$ as $c \uparrow M$, which implies that necessarily $H'(c) \uparrow \infty$ as $c \uparrow M$. To prove the claim, let $0 < \beta < c < M$. Then since $H'(s)$ is increasing

$$H(c) = \int_0^c H'(s) ds \leq \beta H'(\beta) + (c - \beta)H'(c).$$

This implies

$$cH'(c) - H(c) \geq \beta \left[H'(c) - H'(\beta) \right].$$

Letting $c \uparrow M$ and using that $\beta > 0$, the claim is proved, thus completing the proof of the theorem. ■

One of the situations of interest when aleatoric and epistemic variables appear simultaneously is the case where all that is known regarding the epistemic variables are bounds. It turns out that this problem is well posed (i.e., the relevant performance criteria are finite) essentially in those cases where $\Lambda_\infty^1 < \infty$. In fact, when this is the case the value of Λ_∞^1 can be used to establish optimal bounds subject to the constraint on the epistemic variables. For simplicity, we assume that the set A appearing in the statement of the following theorem is bounded and that γ , the nominal distribution for the epistemic variables, is the uniform distribution (in the limit $c \rightarrow \infty$ the precise form of the nominal distribution is not important, and in fact it is only the support of the distribution that matters in the limit). If A is not bounded one can use an appropriate distribution whose support is all of A (e.g., the exponential distribution could be used for $A = [0, \infty)$). The choice of γ as uniform when A is bounded is simplest and also one that is convenient for most uses. A related statement holds for Λ_∞^2 , but it is not as useful (at least in this setting), since $\Lambda_\infty^1 \leq \Lambda_\infty^2$.

Theorem 4 *Suppose that $A \subset \mathcal{Y}$ is bounded and the closure of its interior and that γ is the uniform distribution on A . Assume that F is lower semicontinuous in y for each $x \in \mathcal{X}$ and bounded from below. Define the risk-sensitive functional Λ_c^1 by (4) and let $\Lambda_\infty^1 = \lim_{c \rightarrow \infty} \Lambda_c^1$. Then*

$$\sup_{y \in A} \int_{\mathcal{X}} F(x, y) \mu(dx) = \Lambda_\infty^1.$$

Proof. Since $c\Lambda_c^1$ is convex the limit of Λ_c^1 is well defined, though it might take the value ∞ . We give the proof for the case $\Lambda_\infty^1 < \infty$. The extension to the case $\Lambda_\infty^1 = \infty$ is straightforward.

We first prove that the left side of the last display is bounded above by the right side. Fix any $\bar{y} \in A^\circ$, the interior of A . For small $\varepsilon > 0$ let B_ε be the ball in A about \bar{y} of volume ε , and let M be the volume of A . Let $\theta(dy)$ be

the measure whose density with respect to Lebesgue measure is $1/\varepsilon$ on B_ε , and zero elsewhere. Then

$$\begin{aligned} R(\theta(dy) \parallel \gamma(dy)) &= \int_{B_\varepsilon} \log[(1/\varepsilon)/(1/M)] \theta(dy) \\ &= (1/\varepsilon) \log[(1/\varepsilon)/(1/M)]. \end{aligned}$$

Then (5) gives

$$\int_{\mathcal{X}} \int_{B_\varepsilon} F(x, y) \theta(dy) \mu(dx) \leq \frac{1}{c} (1/\varepsilon) \log[(1/\varepsilon)/(1/M)] + \Lambda_c^1. \quad (8)$$

Since $y \rightarrow F(x, y)$ is lower semicontinuous

$$\liminf_{\varepsilon \rightarrow 0} \int_{B_\varepsilon} F(x, y) \theta(dy) \geq F(x, \bar{y}).$$

Letting first $c \rightarrow \infty$ and then $\varepsilon \rightarrow 0$ in (8), Fatou's lemma gives $\int_{\mathcal{X}} F(x, \bar{y}) \mu(dx) \leq \Lambda_\infty^1$. Since $\bar{y} \in A^\circ$ is arbitrary, the lower semicontinuity and another use of Fatou give the bound for all $y \in A$.

For the reverse inequality, we use the fact that the minimizing measure in the definition of Λ_c^1 exists. In fact,

$$\int_{\mathcal{Y}} \int_{\mathcal{X}} F(x, y) \mu(dx) \theta_c^*(dy) = \frac{1}{c} R(\theta_c^*(dy) \parallel \gamma(dy)) + \Lambda_c^1$$

precisely at θ^* defined by

$$\frac{d\theta_c^*(\cdot)}{d\gamma(\cdot)}(y) = e^{\int_{\mathcal{X}} cF(x, y) \mu(dx)} \Big/ \int_A e^{\int_{\mathcal{X}} cF(x, y) \mu(dx)} \gamma(dy)$$

(see [2, Proposition 1.4.2]). By the non-negativity of relative entropy

$$\begin{aligned} \Lambda_\infty^1 &= \lim_{c \rightarrow \infty} \Lambda_c^1 \\ &\leq \limsup_{c \rightarrow \infty} \int_{\mathcal{Y}} \int_{\mathcal{X}} F(x, y) \mu(dx) \theta_c^*(dy) \\ &\leq \sup_{y \in A} \int_{\mathcal{X}} F(x, y) \mu(dx). \end{aligned}$$

■

4 Examples and computational methods

To indicate how the robust performance bounds might be used in practice, we present some numerical examples, which illustrate the relationship between the different risk sensitive integrals and the relevant techniques and computational issues. We first review the polynomial chaos techniques that are used to compute the risk-sensitive integrals. The reader familiar with this material can skip to the next section.

Remark 5 Although previously the random variables with known and unknown distributions were denoted by X and Y , in the rest of the paper these random variables will be denoted by Z_1 and Z_2 . This is done to free up x and y to be spatial variables in various PDE and related equations that might be used to define the mapping F .

4.1 Review of polynomial chaos methods

This subsection reviews the polynomial chaos method that we use to compute the risk-sensitive integrals. The reader familiar with this material can skip to Section 4.2.

To calculate risk sensitive integrals, one could use Monte Carlo integration, which requires evaluating $F(Z_1, Z_2)$ for many replicas of (Z_1, Z_2) . This can get costly (though not necessarily for the given examples), and so we utilize generalized polynomial chaos (gPC) methods to approximate $F(z_1, z_2)$ and evaluate various statistical properties of $F(Z_1, Z_2)$ (see, for example, [9]). We briefly recall the mathematical framework of polynomial chaos methods, along with a simple example in one dimension.

Let $(\Omega, \mathcal{A}, \mathcal{P})$ be a probability space where Ω is the sample space, \mathcal{A} the associated σ -algebra, and \mathcal{P} the probability measure on \mathcal{A} . Define a random vector $\mathbf{Z}(\omega) \doteq (Z_1(\omega), \dots, Z_N(\omega)) \in \mathbb{R}^N$ on this probability space. Consider a d -dimensional bounded domain $D \subset \mathbb{R}^d$ with boundary ∂D and let $t \in [0, T]$, $T \in (0, \infty)$. We consider random fields $u(t, x; \mathbf{Z}(\omega)) : \bar{D} \times [0, T] \times \Omega \rightarrow \mathbb{R}$ that are defined by requiring that for a.e. $\omega \in \Omega$

$$\mathcal{L}(t, x, u; \mathbf{Z}(\omega)) = f(t, x; \mathbf{Z}(\omega)), \quad (x, t, \omega) \in D \times [0, T] \times \Omega,$$

subject to the boundary and initial conditions

$$\begin{aligned} \mathcal{B}(t, x, u; \mathbf{Z}(\omega)) &= g(t, x; \mathbf{Z}(\omega)), & (x, t, \omega) \in \partial D \times [0, T] \times \Omega, \\ u(t, x; \mathbf{Z}(\omega)) &= u_0(x; \mathbf{Z}(\omega)), & (x, t, \omega) \in D \times \{0\} \times \Omega. \end{aligned}$$

Here $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, \mathcal{L} is a linear or nonlinear differential operator, and \mathcal{B} is a boundary operator. We assume for simplicity that a unique classical sense solution exists to the differential equation and/or boundary conditions. Note that for the first two examples of the last subsection the problem involves only time dependence, and hence there is neither a spatial variable nor a boundary condition.

We assume that $\{Z_i\}_{i=1}^N$ are independent random variables, with support $\{\Gamma_i\}_{i=1}^N$ and with probability density functions $\{\rho_i(z_i)\}_{i=1}^N$, respectively. Hence the joint density is $\rho(\mathbf{z}) = \prod_{i=1}^N \rho_i(z_i)$. Let Ω be the canonical space $\times_{i=1}^N \Gamma_i$, in which case $\mathbf{Z}(\omega) = \omega$ and we identify ω with \mathbf{z} . Thus the equations above become

$$\mathcal{L}(t, x, u; \mathbf{z}) = f(t, x; \mathbf{z}), \quad (x, t, \mathbf{z}) \in D \times [0, T] \times \Omega, \quad (9)$$

subject to the boundary and initial conditions

$$\begin{aligned} \mathcal{B}(t, x, u; \mathbf{z}) &= g(t, x; \mathbf{z}), & (x, t, \mathbf{z}) \in \partial D \times [0, T] \times \Omega, \\ u(t, x; \mathbf{z}) &= u_0(x; \mathbf{z}), & (x, t, \mathbf{z}) \in D \times \{0\} \times \Omega. \end{aligned} \quad (10)$$

We consider approximating the mapping as a function of the stochastic variable, i.e., $\mathbf{z} \rightarrow u(t, x; \mathbf{z})$ at some particular (t, x) , via a finite sum of orthogonal basis functions.

We first define finite dimensional subspaces for $L^2(\Gamma_i)$ according to

$$W_i^{P_i} = \left\{ v : \Gamma_i \rightarrow \mathbb{R} : v \in \text{span} \{ \phi_{i,m}(z_i) \}_{m=0}^{P_i} \right\}, \quad i = 1, \dots, N.$$

Here P_i represents the highest degree of the polynomial basis function, and $\{ \phi_{i,m}(z_i) \}$ are a set of orthonormal polynomials with respect to the weight ρ_i , i.e., for $m \neq n$

$$\int_{\Gamma_i} \phi_{i,m}(z_i) \phi_{i,n}(z_i) \rho_i(z_i) dz_i = 0$$

and

$$\int_{\Gamma_i} \phi_{i,n}^2(z_i) \rho_i(z_i) dz_i = 1.$$

The orthonormal basis representation is determined by the probability density function ρ_i . For example, if the density is uniform or Gaussian, then Legendre or Hermite orthogonal polynomials, respectively, are used. A table of polynomial basis functions and their respective distributions is listed at the end of this section (also see [6]). A finite dimensional subspace for $L^2(\Gamma)$, where $\Gamma \doteq \Gamma_1 \times \dots \times \Gamma_N = \Omega$, can either be defined as

$$W_N^P = \bigotimes_{|\mathbf{P}| \leq P} W_i^{P_i}$$

where the tensor product is over all combinations of the multi-index $\mathbf{P} = (P_1, \dots, P_N) \in \mathbb{N}_0^N$ with $|\mathbf{P}| = \sum_{i=1}^N P_i \leq P$, or

$$\tilde{W}_N^P = \bigotimes_{i=1}^N W_i^P.$$

Thus, W_N^P is the space of N -variate orthogonal polynomials of total degree at most P , whereas \tilde{W}_N^P is the full tensor product of the one-dimensional polynomial spaces with each highest degree P . Note that $\dim(W_N^P) = \binom{N+P}{P}$ and $\dim(\tilde{W}_N^P) = (P+1)^N$, and that for large N , $\binom{N+P}{P} \ll (P+1)^N$. Since our examples only consider $N = 2$, we will use the full tensor product space, \tilde{W}_N^P . Let $\Phi_j(\mathbf{z}), j = 1, \dots, M$ denote the elements of \tilde{W}_N^P , where $M = \dim(\tilde{W}_N^P)$.

There are two standard methods for constructing gPC approximations, referred to as the stochastic Galerkin method (see, e.g., [6]) and the stochastic collocation method (see, e.g., [7]). Since the calculations presented below use only the collocation method, we restrict discussion to this method. With stochastic collocation, we first consider an approximation to the solution $u(t, x; \mathbf{z})$ in terms of a Lagrange interpolant of the form

$$v(t, x; \mathbf{z}) = \sum_{k=1}^{N_p} v(t, x; \mathbf{z}_k) \mathbf{L}_k(\mathbf{z}), \quad (11)$$

where $\mathbf{L}_k(\mathbf{z})$ is a Lagrange polynomial of degree $\dim(W_N^P)$, satisfies $\mathbf{L}_k(\mathbf{z}_l) = \delta_{kl}$, and $\{\mathbf{z}_k\}_{k=1}^{N_p}$ are a set of prescribed nodes in the N -dimensional space Γ . We require that the residuals

$$\begin{aligned} R(t, x; \mathbf{z}) &\doteq \mathcal{L}(t, x, v; \mathbf{z}) - f(t, x; \mathbf{z}) \\ R^{BC}(t, x; \mathbf{z}) &\doteq \mathcal{B}(t, x, v; \mathbf{z}) - g(t, x; \mathbf{z}) \\ R^{IC}(x; \mathbf{z}) &\doteq v(0, x; \mathbf{z}) - u_0(x; \mathbf{z}) \end{aligned}$$

vanish at the collocation points $\{\mathbf{z}_k\}_{k=1}^{N_p}$, i.e.,

$$\begin{aligned} \mathcal{L}(t, x, v(t, x; \mathbf{z}_k); \mathbf{z}_k) - f(t, x; \mathbf{z}_k) &= 0, \text{ for } k = 1, \dots, N_p \\ \mathcal{B}(t, x, v(t, x; \mathbf{z}_k); \mathbf{z}_k) - g(t, x; \mathbf{z}_k) &= 0, \text{ for } k = 1, \dots, N_p \\ v(0, x; \mathbf{z}_k) - u_0(x; \mathbf{z}_k) &= 0, \text{ for } k = 1, \dots, N_p. \end{aligned}$$

Thus, the stochastic collocation method produces a set of N_p decoupled equations, where each value of $v(t, x; \mathbf{z}_k)$ coincides with the exact solution $u(t, x; \mathbf{z})$ for the given \mathbf{z}_k , since both satisfy the same equations.

Once the values of $\{v(t, x; \mathbf{z}_k)\}_{k=1}^{N_p}$ have been determined at the collocation points, one can construct a gPC approximation in the orthogonal basis representation in the form

$$v(t, x; \mathbf{z}) = \sum_{j=1}^M \hat{v}_j(t, x) \Phi_j(\mathbf{z}), \quad (12)$$

which, under certain conditions, will be equivalent to the Lagrange basis formulation. The coefficients $\{\hat{v}_j\}_{j=1}^M$ can be determined by inverting a Vandermonde matrix, where invertibility is dependent on the choice of collocation points. If one chooses quadrature or cubature¹ points for the collocation points, then invertibility is guaranteed. Furthermore, if the cubature rule is exact up to polynomials of degree $2M$, inversion of the Vandermonde matrix is not necessary², and the coefficients $\{\hat{v}_j\}_{j=1}^M$ can be computed by evaluating

$$\hat{v}_j(t, x) = \sum_{k=1}^{N_p} v(t, x; \mathbf{z}_k) \Phi_j(\mathbf{z}_k) w_k, \quad (13)$$

where $\{w_k\}_{k=1}^{N_p}$ are the respective weights according to the choice of quadrature or cubature points.

Statistical information is easy to obtain from the form (12). For example,

$$\mathbb{E}_{\mathbf{z}}[v(t, x; \mathbf{z})] = \int_{\Gamma} \left(\sum_{j=1}^M \hat{v}_j(t, x) \Phi_j(\mathbf{z}) \right) \rho(\mathbf{z}) d\mathbf{z} = \hat{v}_1(t, x)$$

¹cubature points just refers to quadratures in more than one dimension.

²This is especially useful if the Vandermonde matrix is ill-conditioned.

and

$$\mathbb{E}_{\mathbf{z}}[(v(t, x; \mathbf{z}))^2] = \int_{\Gamma} \left(\sum_{j=1}^M \hat{v}_j(t, x) \Phi_j(\mathbf{z}) \right)^2 \rho(\mathbf{z}) d\mathbf{z} = \sum_{j=1}^M \hat{v}_j(t, x)^2$$

as a consequence of the orthonormality. Other statistical information, such as higher moments and sensitivity coefficients, are also easily calculated (see [6]). The ease with which one can calculate such moments explains why orthogonal representation (12) is preferred over the Lagrange form (11).

To illustrate the approach, we look at a simplified version of Example 1. Let $g(z_2) = u_0$, so that the only uncertainty is in the decay rate k , which depends on the random variable Z . The simplified problem becomes

$$\frac{d}{dt}u(t) = -k(z)u(t), \quad u(0) = u_0,$$

where $u(t; z)$ denotes the solution parameterized by z , with z a point in the range of Z . Requiring that the residuals vanish at the collocation points $\{z_n\}_{n=1}^{N_p}$ gives

$$R(t; z_n) = \frac{d}{dt}v(t; z_n) + k(z_n)v(t; z_n) = 0$$

and

$$R^{IC}(z_n) = v(0; z_n) - u_0 = 0.$$

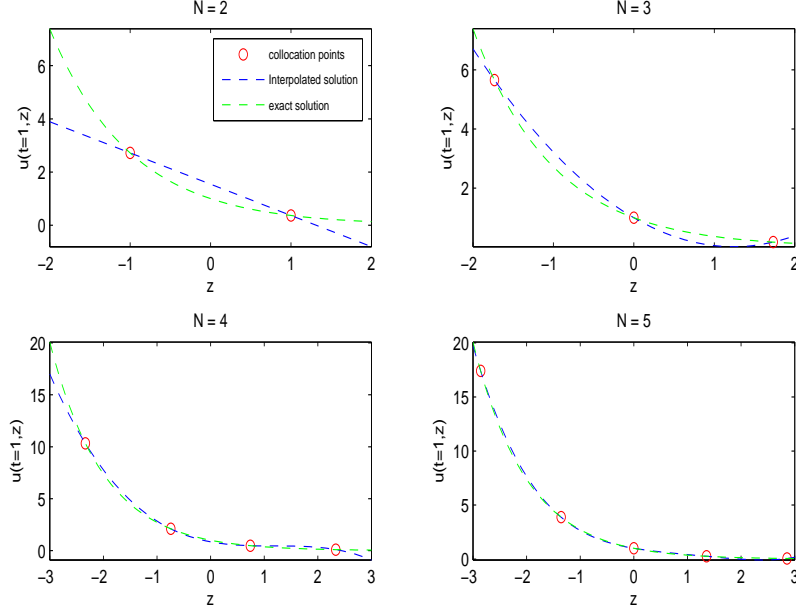
Note that there are indeed N_p decoupled equations, which are solved independently for each n . Here the solutions are exact:

$$v(t; z_n) = u_0 e^{-k(z_n)t}, \text{ for } 1, \dots, N_p.$$

One can then obtain the approximation (12) with respect to the orthogonal basis, assuming that the appropriate collocation points and weights have been chosen.

We illustrate this example with $k(z) = z$ and $Z \sim \mathcal{N}(\mu = 0, \sigma^2 = 1)$. Here the $\{\Phi_j\}_{j=1}^M$ are the Hermite polynomials, which are orthonormal with respect to the weighting function $\rho(z) = e^{-(z-\mu)^2/2\sigma^2}/\sigma\sqrt{2\pi}$ with support $\Gamma = (-\infty, +\infty)$. Figure 2 shows the stochastic collocation approximation $v(1, z)$ at $t = 1$ as a function of z , with $N_p = 2, 3, 4, 5$, versus the exact solution $u_{exact}(1, z) = e^{-z}$. Note that the collocation approximation is an interpolation of the exact solution at N_p points.

Figure 3 shows the relative error between the exact and approximated mean and variance in a semilog plot. The linear decrease in relative error on the semilog plot implies an exponential decay in the error. Errors become a constant at the limits of the accuracy of the numerical scheme.

Figure 2: Interpolated solutions versus the exact solution at $t = 1$.

4.2 Numerical examples and interpretation

We now discuss numerical calculation of the risk-sensitive integrals. For convenience, we recall the three types of risk-sensitive integrals introduced in Section 3.1, the ordinary risk sensitive cost Λ_c and the two hybrid integrals Λ_c^1 and Λ_c^2 :

$$\begin{aligned}\Lambda_c &= \frac{1}{c} \log \int_{\mathcal{Z}_1} \int_{\mathcal{Z}_2} e^{cF(z_1, z_2)} \gamma(dz_2) \mu(dz_1), \\ \Lambda_c^1 &= \frac{1}{c} \log \int_{\mathcal{Z}_2} e^{\int_{\mathcal{Z}_1} cF(z_1, z_2) \mu(dz_1)} \gamma(dz_2), \\ \Lambda_c^2 &= \frac{1}{c} \int_{\mathcal{Z}_1} \left[\log \int_{\mathcal{Z}_2} e^{cF(z_1, z_2)} \gamma(dz_2) \right] \mu(dz_1),\end{aligned}\tag{14}$$

where $\mathcal{Z}_1, \mathcal{Z}_2$ are the range spaces for the random variables Z_1, Z_2 , respectively. Recall the relationships $\Lambda_c^1 \leq \Lambda_c^2 \leq \Lambda_c$. It is assumed that the distribution of Z_1 is known, but this is not true for Z_2 .

In the stochastic collocation approach, we replace $F(z_1, z_2)$ in (14) with a polynomial interpolant, $\tilde{F}(z_1, z_2)$, determined by the full tensor product of the




Figure 3: Relative errors.

one dimensional collocation points, $\{z_{1,i}\}_{i=1}^{N_{1,q}} \otimes \{z_{2,i}\}_{i=1}^{N_{2,q}}$ in $\mathcal{Z}_1 \times \mathcal{Z}_2$. Thus

$$\tilde{F}(z_1, z_2) = \sum_{j=1}^M \hat{F}_j(t, x) \Phi_j(z_1, z_2),$$

where $M = N_{1,q} \cdot N_{2,q}$, Φ_j is the full tensor product basis, and $\hat{F}_j(t, x)$ is computed using (13). From here, one can efficiently compute the risk sensitive integrals (14) via either Monte Carlo sampling or quadrature. For both methods, once the samples or quadrature points have been chosen, one can calculate the risk sensitive integrals for different values of c without resampling or choosing different quadrature points. Hence, it is computationally inexpensive to compute the risk sensitive integral for different values of c . Note that for higher dimensions, a sparse grid is preferred and, in most cases, necessary for these types of computations. It is typical to use the Smolyak algorithm to generate these sparse grids, which are based on a one dimensional quadrature rule (for example, see Section 4.1.2 in [5] for a more detailed explanation and further references).

In this paper, we opt for the numerical quadrature approach. Since $\tilde{F}(z_1, z_2)$ can be evaluated very quickly, we can compute the risk sensitive integrals with high accuracy and efficiency using a high order quadrature rule. The number of quadrature points chosen to compute the risk sensitive integrals need not be the same as the number of collocation points used in computing the coefficients of the gPC expansion, though the type of quadrature points is the same. For example, if we were to use Legendre-Gauss quadrature points to evaluate

the gPC coefficients, we would again use Legendre-Gauss quadrature points to calculate the risk sensitive integrals. In the examples presented below a much larger number of quadrature points are used to compute the risk sensitive integrals. To avoid confusion, the set of quadrature points used to evaluate the risk sensitive integrals is denoted by $\{\tilde{z}_{1,i}\}_{i=1}^{\tilde{N}_{1,q}} \otimes \{\tilde{z}_{2,i}\}_{i=1}^{\tilde{N}_{2,q}}$, i.e., the full tensor product of one dimensional quadrature points. In the examples, $\tilde{N}_{1,q}$ and $\tilde{N}_{2,q}$ equal 2^8 , while $N_{1,q}$ and $N_{2,q}$ are 8 for smooth F and 12 when F is an indicator functions.

We also implemented the Monte Carlo approach, but observed that numerical quadrature gave a much better approximation for the same computational effort, at least for our examples. It should also be noted that a variation on traditional Monte Carlo is needed for the hybrid forms of the risk sensitive integrals. For the original risk-sensitive cost function (3), the standard approach is to use many independent samples of the pair (Z_1, Z_2) . This produces an estimate which has a mean square error of $\mathcal{O}(N^{-1/2})$, where N is the total number of Monte Carlo samples. However, for the first hybrid form (4), one must approximate the inner integral $\int_{\mathcal{X}} cF(z_1, z_2) \mu(dz_1)$ for every fixed sample of Z_2 . Thus one has to simulate more Z_1 samples than Z_2 samples. Similarly, to compute the second form of the hybrid (6), one must approximate the inner integral $\int_{\mathcal{Y}} e^{cF(z_1, z_2)} \gamma(dz_2)$ for every fixed Z_1 . In particular, for both hybrid forms one does not simply choose independent samples of (Z_1, Z_2) .

4.2.1 Independent aleatoric and epistemic uncertainties

When the aleatoric and nominal epistemic variables are independent, the numerical quadrature approach would evaluate the risk sensitive integrals via the formulas

$$\begin{aligned}\Lambda_c &\approx \frac{1}{c} \log \left(\sum_{i=1}^{\tilde{N}_{1,q}} \sum_{j=1}^{\tilde{N}_{2,q}} e^{cF(\tilde{z}_{1,i}, \tilde{z}_{2,j})} \tilde{w}_{1,i} \tilde{w}_{2,j} \right) \\ \Lambda_c^1 &\approx \frac{1}{c} \log \left(\sum_{j=1}^{\tilde{N}_{2,q}} e^{\left(\sum_{i=1}^{\tilde{N}_{1,q}} cF(\tilde{z}_{1,i}, \tilde{z}_{2,j}) \tilde{w}_{1,i} \right)} \tilde{w}_{2,j} \right) \\ \Lambda_c^2 &\approx \frac{1}{c} \sum_{i=1}^{\tilde{N}_{1,q}} \left(\log \sum_{j=1}^{\tilde{N}_{2,q}} e^{cF(\tilde{z}_{1,i}, \tilde{z}_{2,j})} \tilde{w}_{2,j} \right) \tilde{w}_{1,i}.\end{aligned}$$

where $\{\tilde{z}_{1,i}, \tilde{w}_{1,i}\}_{i=1}^{\tilde{N}_{1,q}}, \{\tilde{z}_{2,i}, \tilde{w}_{2,i}\}_{i=1}^{\tilde{N}_{2,q}}$ are the pairs of quadrature points and weights chosen based on the underlying distribution.

The first example is taken from [8] and involves a simple ODE. The example is convenient because analytic solutions to the risk-sensitive integrals can be obtained, and then compared with the numerical approximations described in the next section. The second example is a random oscillator, where the stochas-

tic parameters are the spring and dampening coefficients. The third example involves a heat equation with random heat capacity and thermal conductivity.

Example 1 We consider $F(Z_1, Z_2)$ defined in terms of the solution of the ODE

$$\frac{d}{dt}u(t) = -k(z_1)u(t), \quad u(0) = g(z_2),$$

where k and g are known functions. Letting $u(t; z_1, z_2)$ denote the solution parameterized by (z_1, z_2) , set $F(Z_1, Z_2) = h(u(t; Z_1, Z_2))$, where typical choices of h are $h(u) = u$, $h(u) = u^2$, $h(u) = \mathbf{1}\{a \leq u \leq b\}$.

Here we take $k(z_1) = z_1$, $g(z_2) = z_2$, let $Z_1 \sim U[0, 1]$, $Z_2 \sim U[0, 1]$, and compute risk sensitive integrals for $F(z_1, z_2) = [u(1; z_1, z_2)]^2$, and $F(z_1, z_2) = \mathbf{1}\{.8 \leq u(1; z_1, z_2) \leq 1\}$. Hence the goal is to obtain robust bounds on the second moment and the probability that the solution to the stochastic ODE falls within the interval $[.8, 1]$ at time $t = 1$, respectively. Note that the indicator function is not smooth, and thus more collocation points are required to obtain an accurate approximation. We choose a Legendre polynomial basis for both Z_1 and Z_2 and use the Legendre-Gauss quadrature points and weights for each dimension.

Figures 4 and 5 show the approximations as a function of c for $F(u) = u^2$ and $F(u) = \mathbf{1}\{.8 \leq u \leq 1\}$, respectively. As expected Λ_c^1 gives the best bounds, while the original form Λ_c gives the worst upper bounds of the three.




Figure 4: Example 1, $\{\Lambda_c^i\}_{i=1}^2$ for $F(u) = u^2$.

The plots also suggest what happens as $c \rightarrow \infty$. The measure Λ_c is expected to yield robust bounds if one is uncertain regarding the distributions of both random variables Z_1 and Z_2 , and $\lim_{c \rightarrow \infty} \Lambda_c$ represents the tightest upper bound




Figure 5: Example 1, $\Lambda_c, \{\Lambda_c^i\}_{i=1}^2$ for $F(u) = \mathbf{1}\{.8 \leq u \leq 1\}$.

on performance if we know nothing about these random variables except for their support. In this case, the limit in Figure 5 appears to be 1, which means that for some choice of z_1 and z_2 , $u(1; z_1, z_2) \in [.8, 1]$. Hence without more information on the distributions of Z_1, Z_2 we do not obtain information other than this from this performance bound. On the other hand, $\lim_{c \rightarrow \infty} \Lambda_c^i$ for $i = 1, 2$ is strictly less than 1 and thus by Theorem 4 gives a meaningful performance bound when the only information available about Z_2 is its support. Note also that while Λ_c^1 and Λ_c^2 seem close for $F(u) = u^2$, they differ considerably for $F(u) = \mathbf{1}\{.8 \leq u \leq 1\}$.

In Figures 6 and 7 we display the relative error $|\Lambda_c^i - \Lambda_{c,exact}^i| / |\Lambda_{c,exact}^i|$ between the three risk sensitive integrals and the exact solution, which we were able to calculate using partially analytical solutions. The relative error is fairly small for $F(u) = u^2$, but significantly larger for $F(u) = \mathbf{1}\{.8 \leq u \leq 1\}$, even though more collocation points were used to determine the gPC coefficients.

Next we consider the problem of computing the value of c which minimizes $c \rightarrow \frac{1}{c}B + \Lambda_c^i$ for $i = 0, 1, 2$.³ Here B is a maximum relative entropy distance that will be allowed between the “true” distribution, $\theta(dz_2)$, and the nominal distribution, $\gamma(dz_2)$, and the minimization produces the tightest possible bounds. For the example, suppose that the family of distributions for which bounds are to be valid includes $\theta(dz_2) \sim \text{beta}(\alpha = 1.5, \beta = 1.5)$, which implies $B \approx .0484$. In this example $B \leq f(c) \doteq c \frac{d}{dc} H^i(c) - H^i(c)$ for some c , where $H^i(c) \doteq c \Lambda_c^i$. As shown in Proposition 3, this implies there is a unique local minimum for $c \in (0, \infty)$. See Figures 8 and 9. Alternatively, if $B > f(c)$ for all c , then Proposition 3 implies the minimum is achieved in the limit as $c \rightarrow \infty$. From

³Note that the risk sensitive integrals are infinite for $c = 0$. Thus in the numerical calculations, we actually compute c starting from .01.




Figure 6: Example 1, relative error for Λ_c , $\{\Lambda_c^i\}_{i=1}^2$ with $F(u) = u^2$.




Figure 7: Example 1, relative error for Λ_c , $\{\Lambda_c^i\}_{i=1}^2$ with $F(u) = \mathbf{1}\{.8 \leq u \leq 1\}$.

Figure 9 we find a minimum for $\frac{1}{c}B + \Lambda_c^1$ of approximately 0.04 at $c \approx 5.12$. Thus for all distributions on Z_2 whose relative entropy distance to $U[0, 1]$ is less than 0.0484, the probability that the solution to the random ODE falls between .8 and 1 at time 1 is less than 0.04.

Note that the minimum is at different values of c for the three different integrals, but the minimum is always the smallest for the first form of the hybrid. In the case when the minimum occurs at some $c < \infty$, the minimum is easily calculated, since this is a one-dimensional unconstrained minimization




Figure 8: Example 1, $\frac{1}{c}(B + \Lambda_c)$, $\{\frac{1}{c}(B + \Lambda_c^i)\}_{i=1}^2$ with $F(u) = u^2$.




Figure 9: Example 1, $\frac{1}{c}(B + \Lambda_c)$, $\{\frac{1}{c}(B + \Lambda_c^i)\}_{i=1}^2$ with $F(u) = \mathbf{1}\{.8 \leq u \leq 1\}$.

problem (for example, one can use a golden method search algorithm, such as MATLAB's `fminbnd` function). In the case when the minimum occurs in the limit as $c \rightarrow \infty$, in order to find the minimum, one can iterate Λ_c^i until the successive iterations differ by less than a prescribed tolerance.

The second example is a random oscillator, where the stochastic parameters are the spring and dampening coefficients.

Example 2 We consider $F(Z_1, Z_2)$ defined in terms of the solution of the ODE

$$\frac{d^2}{dt^2}u(t) + \gamma(z_2)\frac{d}{dt}u(t) + k(z_1)u(t) = f(t), \quad u(0) = u_0, \quad u'(0) = u'_0,$$

where k and γ are known functions, which represent the spring and dampening coefficients, respectively. The outcome of interest is whether or not the position of the random oscillator falls within a specified range, $[a, b]$, at a specified time, $t_{critical}$. Letting $u(t; z_1, z_2)$ denote the solution parameterized by (z_1, z_2) , set $F(Z_1, Z_2) = \mathbf{1}\{a \leq u(t_{critical}; Z_1, Z_2) \leq b\}$.

Here we consider the random oscillator with $k(z_1) = 4z_1$, $\gamma(z_2) = z_2/5$, and $f(t) = 10 \cos(10t) + 3$, and choose $Z_1 \sim \text{beta}(\alpha = 5, \beta = 5)$, $Z_2 \sim \text{beta}(\alpha = 5, \beta = 5)$, $t_{critical} = 4$, and $[a, b] = (-\infty, 2]$. Here we use the Gauss-Jacobi quadrature points and weights. In this example we are concerned with the position of the oscillator at a critical time. Figure 10 plots the risk sensitive integrals for Example 2 with $F(u) = \mathbf{1}\{u \leq 2\}$. Figure 11 depicts $c \rightarrow \frac{1}{c}(B + \Lambda_c^i)$, where $B = R(\theta(dz_2) \|\gamma(dz_2)) \approx .0587$, with $\theta(dz_2) \sim \text{beta}(\alpha = 10, \beta = 10)$.




Figure 10: Example 2, $\Lambda_c, \{\Lambda_c^i\}_{i=1}^2$ for $F(u) = \mathbf{1}\{u \leq 2\}$.

Finally, the third example involves a heat equation with random heat capacity and thermal conductivity.

Example 3 We consider $F(Z_1, Z_2)$ defined in terms of the solution of the PDE

$$\frac{d}{dt}u(t, x) = \frac{k(u; z_1)}{m(z_2)} \frac{d^2}{dx^2}u(t, x), \quad u(0, x) = u_0(x)$$

on $x \in (0, L)$ with boundary conditions

$$-k(u; z_1) \frac{d}{dx}u(t, 0) = q, \quad \frac{d}{dx}u(t, L) = 0.$$




Figure 11: Example 2, $\frac{1}{c}(B + \Lambda_c), \{\frac{1}{c}(B + \Lambda_c^i)\}_{i=1}^2$ with $F(u) = \mathbf{1}\{u \leq 2\}$.



Here k and m are the thermal conductivity and heat capacity, respectively. Note that in this example, the randomness affects the diffusivity and the boundary conditions. We are interested in the probability that the material exceeds a particular temperature, $T_{critical}$, at the time t_{final} and at some point $x^* \in [0, L]$. Letting $u(t; z_1, z_2)$ denote the solution parameterized by (z_1, z_2) , set $F(Z_1, Z_2) = \mathbf{1}\{u(t_{final}, x^*; Z_1, Z_2) \geq T_{critical}\}$.

Here we consider the random heat equation with (nonlinear) Neumann boundary conditions. We use $k(u, z_1) = z_1 + (1.5 \times 10^{-7})u$, $m(z_2) = (10^{-6})z_2$, $q = .35$, $L = 1.90$, $u_0(x) = 25$, and $T_{critical} = 980$, $x^* = 0$, $t_{final} = 1000$, and set $F(Z_1, Z_2) = \mathbf{1}\{u(t_{final}, x^*; Z_1, Z_2) \geq T_{critical}\}$. For the random variables, we introduce two independent random variables $\tilde{Z}_1 \sim \text{beta}(\alpha = 5, \beta = 5)$, $\tilde{Z}_2 \sim \text{beta}(\alpha = 5, \beta = 5)$ and let $Z_1 = (2 \times 10^{-3})\tilde{Z}_1 + 3 \times 10^{-3}$, $Z_2 = (0.11)\tilde{Z}_2 + 0.30$. Figures 12 and 13 plot the risk sensitive integrals as a function of c and illustrate $c \rightarrow \frac{1}{c}(B + \Lambda_c^i)$, where $B = R(\theta(dz_2) \|\gamma(dz_2)) \approx .0587$. and with $\theta(dz_2) \sim \text{beta}(\alpha = 10, \beta = 10)$. Again, we use the Gauss-Jacobi quadrature points and weights.

4.2.2 Dependent aleatoric and epistemic uncertainties

In this section we consider problems where the distribution of Z_1 can depend on the value of Z_2 , or vice versa. In the case when Z_1 depends on Z_2 , we use the following inequality, which holds for the alternative of the first hybrid form (see the discussion in Section 3.1):

$$\int_{Z_2} \int_{Z_1} F(z_1, z_2) \mu(dz_1 | z_2) \theta(dz_2) \leq \frac{1}{c} R(\theta(dz_2) \|\gamma(dz_2)) + \bar{\Lambda}_c^1,$$

Figure 12: Example 3, risk sensitive integrals as a function of c .Figure 13: Example 3, plot of $\frac{1}{c}(B + \Lambda_c)$, $\{\frac{1}{c}(B + \Lambda_c^i)\}_{i=1}^2$.

where

$$\bar{\Lambda}_c^1 = \frac{1}{c} \log \int_{\mathcal{Z}_2} \exp \left(\int_{\mathcal{Z}_1} cF(z_1, z_2) \mu(dz_1|z_2) \right) \gamma(dz_2).$$

Recall that the relative entropy term $\frac{1}{c}R(\theta(dz_2) \parallel \gamma(dz_2))$ represents our lack of knowledge about the distribution of Z_2 , i.e., the epistemic uncertainty where we want robustness, whereas Z_1 represents aleatoric uncertainty. A question that now arises is whether or not it is still possible to use gPC methods to evaluate

this integral. Consider the example with $Z_2 \sim U[0, 1]$ and $Z_1 \sim N(Z_2, 1)$, so that Z_1 is Gaussian with variance 1 and mean Z_2 . Certainly Z_1 and Z_2 are not independent random variables. However, we can introduce an auxiliary normal random variable $Z \sim \mathcal{N}(0, 1)$, and write $Z_1 \doteq Z + Z_2$, where Z_2 and Z are independent random variables. Now we consider $G(Z, Z_2)$ instead of $F(Z_1, Z_2)$, and note that $G(Z, Z_2) = F(Z + Z_2, Z_2)$. Furthermore, $\bar{\Lambda}_c^1$ can be written as

$$\bar{\Lambda}_c^1 = \frac{1}{c} \log \int_{Z_2} \exp \left(\int_Z cG(z, z_2) \mu(dz) \right) \gamma(dz_2).$$

Hence, $\bar{\Lambda}_c^1$ can be calculated just like Λ_c^1 via a gPC approximation. Similar results hold for the case when Z_2 depends on Z_1 , but the epistemic uncertainty is in Z_2 .

In general, as long as we can find a smooth, invertible mapping from (Z_1, Z_2) to an independent pair of random variables, one can evaluate the integrals as before. Figures 14 and 15 show the performance of risk sensitive integrals for dependent random variables as above for Example 1 and with $F(u) = \mathbf{1}\{1/2 \leq u \leq 1\}$.




Figure 14: Example 1, $\bar{\Lambda}_c^1$ for $F(u) = \mathbf{1}\{1/2 \leq u \leq 1\}$.

4.3 Discussion and Conclusion

In a recent paper by Xiu et. al. [8], an approach to dealing with epistemic uncertainty via polynomial chaos methods was developed. In [8], epistemic uncertainty meant that the true distribution of the underlying stochastic parameters was not known precisely. Here we compare our method of handling this type of epistemic uncertainty to theirs.




Figure 15: Example $1, \frac{1}{c}(B + \bar{\Lambda}_c^1)$ $F(u) = \mathbf{1}\{1/2 \leq u \leq 1\}$, and $\theta \sim \mathcal{N}(.8, 1)$.

The approach taken in [8] is based on stochastic collocation approximation. Specifically, the differential equation of interest was solved on a generic collocation grid on the sample space of the random variables. This is in contrast to standard uses of stochastic collocation for uncertainty quantification, where the collocation grid is chosen in accordance with the (assumed known) underlying distribution. Of course the reason for the particular choice of basis polynomials and quadrature points is to obtain optimal (spectral) convergence with respect to the L^2 weighted error. In the approach of [8], regardless of what is the “true” (and assumed unknown) distribution, one chooses grid points with respect to the L^2 weighted norm with constant weight (e.g., Legendre-Gauss quadrature weights), even if the original space is unbounded. One then solves the differential equation at these collocation points and constructs the interpolating polynomial of the approximation. Hence, the gPC approximation converges optimally in the L^2 norm only if the underlying distribution is uniform. For non-uniform random variables the approximation will still converge point-wise as the number of collocation points increases (see Section 5 in [8]).

Once the interpolated approximation is obtained, one can determine the mean, variance, probabilities, and so on with respect to various underlying distributions on the parameters via Monte Carlo or other methods. The benefits are still large, in that one need not evaluate the differential equation for every Monte Carlo sample, but instead evaluates a much simpler approximation via a finite sum of polynomial basis functions. However, information can be obtained only by fixing a choice for the underlying distribution. Also, depending on which underlying distribution is used, the errors in the function approximation will have a greater or smaller effect.

In this paper we have developed an approach in which performance bounds

can be calculated for the mean, variance, probabilities, and so on with respect to all distributions that are within a particular relative entropy distance from the nominal distribution, and the bounds are optimal for that class of distributions. In addition to providing performance bounds on a class of distributions defined by their relative entropy distance, the limit $c \rightarrow \infty$ is particularly interesting in that it gives tight bounds on epistemic uncertainties where one assumes nothing about the underlying distribution of certain variables except their support. In addition, the method one can efficiently handle both aleatoric and epistemic uncertainty simultaneously. More precisely, by computing certain hybrid forms of a risk-sensitive expectation we obtain tight performance bounds that distinguish between variables with known distribution and variables with unknown distribution or other forms of epistemic uncertainty. In particular, the scenarios to which the method applies include (i) aleatoric with known distribution; (ii) aleatoric with partly known distribution (mingled aleatoric and epistemic); (iii) epistemic for which one is willing to model by a family of aleatoric uncertainties, and (iv) epistemic where one is only willing to place bounds on the uncertainties. For all the examples we have considered (including those presented in the paper), approximation of the required risk-sensitive integrals can be done via gPC methods using roughly the same computational effort as would be required to compute the corresponding ordinary performance measures using the nominal probability distributions.

A Distribution for Polynomial Basis

Here we list some of the most common distributions, both continuous and discrete, and their corresponding gPC basis representations (see [6]).

Distribution	gPC polynomial basis	
Gaussian	Hermite	
Gamma	Laguerre	
Beta	Jacobi	
Uniform	Legendre	(15)
Poisson	Charlier	
Binomial	Krawtchouk	
Negative Binomial	Meixner	
Hypergeometric	Hahn	

B Relative Entropy Formulas

In this section we evaluate relative entropies for some of the most common types of distributions listed in Table 15. Since most of the distributions listed in this

table fall within the exponential family, we start with the general form

$$f_X(x; \boldsymbol{\theta}) = h(x) \exp \left(\sum_{i=1}^s \eta_i(\boldsymbol{\theta}) T_i(x) - A(\eta_1(\boldsymbol{\theta}), \dots, \eta_s(\boldsymbol{\theta})) \right), \quad (16)$$

where $T_i(x)$, $h(x)$, $\eta_i(\boldsymbol{\theta})$, and $A(\boldsymbol{\theta})$ are known functions, $\boldsymbol{\theta} = (\theta_1, \dots, \theta_d)^T$ is the parameter of the family.

Example 4 For Gaussian distributions $\boldsymbol{\theta} = (\mu, \sigma)^T$, where μ is the mean and σ is the variance. This is put in the general form by setting

$$\boldsymbol{\theta} = (\mu, \sigma)^T, \quad \boldsymbol{\eta} = \left(\frac{\mu}{\sigma^2}, \frac{-1}{2\sigma^2} \right)^T, \quad h(x) = \frac{1}{\sqrt{2\pi}}, \quad T(x) = (x, x^2)^T,$$

$$A(\boldsymbol{\eta}(\boldsymbol{\theta})) = \frac{\mu^2}{2\sigma^2} + \ln |\sigma| = -\frac{\eta_1^2}{4\eta_2} + \frac{1}{2} \ln \left| \frac{1}{2\eta_2} \right|,$$

producing

$$f_X(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/2\sigma^2}.$$

We now state the relative entropy formula for distributions which have a probability density function given by the form of exponential family in (16). The relative entropy between a probability measure P with density $p(x)$ and another probability measure Q with density $q(x)$ is given by

$$R(P \| Q) = \int_{\Gamma} p(x) \ln \frac{p(x)}{q(x)} dx \quad (17)$$

whenever P is absolutely continuous with respect to Q , where Γ is the support of p and q . In all other cases it is defined to be ∞ .

A simplified version of the relative entropy formula for distributions from the same exponential family type is available. Suppose that $q(x)$ and $p(x)$ have the same h, η_i, T_i, A , but with different parameters $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ (note that $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are vectors). If we associate $q(x)$ with $\boldsymbol{\theta}_1$ and $p(x)$ with $\boldsymbol{\theta}_2$, then it can be shown that the relative entropy formula reduces to

$$R(P \| Q) = \sum_{i=1}^s (\eta_i(\boldsymbol{\theta}_2) - \eta_i(\boldsymbol{\theta}_1)) \int_{\Gamma} T_i(x) p(x) dx - A(\boldsymbol{\eta}(\boldsymbol{\theta}_2)) + A(\boldsymbol{\eta}(\boldsymbol{\theta}_1)). \quad (18)$$

Gaussian. If $Q = N(\mu_1, \sigma_1)$ and $P = N(\mu_2, \sigma_2)$, then using the fact that $\int T_1(x) p(x) dx = \mu_2$ and $\int T_2(x) p(x) dx = \mu_2^2 + \sigma_2^2$, we get

$$R(P \| Q) = \frac{1}{2\sigma_1^2} ((\mu_1^2 - \mu_2^2) + (\sigma_2 - \sigma_1)) + \ln \frac{\sigma_1}{\sigma_2}. \quad (19)$$

Beta. For the beta distribution, the parameters in the exponential family are given by

$$\eta = (\alpha - 1, \beta - 1)^T, \quad T = (\ln(x), \ln(1 - x))^T, \quad h(x) = 1, \quad A = \ln \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)}$$

where

$$f_X(x; \alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)} x^{\alpha-1} (1 - x)^{\beta-1}, \quad 0 < x < 1, \quad \alpha, \beta > 1.$$

Another form of the beta distribution, which is more closely related to the Jacobi polynomials, is given by

$$\tilde{f}_X(k; \tilde{\alpha}, \tilde{\beta}) = \frac{(1 - x)^{\tilde{\alpha}} (1 + x)^{\tilde{\beta}}}{2^{\tilde{\alpha} + \tilde{\beta} + 1} B(\tilde{\alpha} + 1, \tilde{\beta} + 1)}, \quad -1 < x < 1, \quad \tilde{\alpha}, \tilde{\beta} > -1,$$

where $B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$ is the beta function. Note that the relation between the two is simply a rescaling and a variable substitution given by $\tilde{\alpha} = \beta - 1, \tilde{\beta} = \alpha - 1$. We will use the primary form to determine the relative entropy formula. Letting $P = \text{beta}(\alpha_2, \beta_2)$ and $Q = \text{beta}(\alpha_1, \beta_1)$,

$$\int_0^1 T_1(x) p(x) dx = \psi(\alpha_2) - \psi(\alpha_2 + \beta_2), \quad \int_{\Gamma} T_2(x) p(x) dx = \psi(\beta_2) - \psi(\alpha_2 + \beta_2),$$

where $\psi(x)$ is known as the digamma function, or the zeroth order of the polygamma function. Then,

$$\begin{aligned} R(P \| Q) &= (\alpha_2 - \alpha_1)(\psi(\alpha_2) - \psi(\alpha_2 + \beta_2)) \\ &\quad + (\beta_2 - \beta_1)(\psi(\beta_2) - \psi(\alpha_2 + \beta_2)) + \ln \frac{B(\alpha_1, \beta_1)}{B(\alpha_2, \beta_2)}. \end{aligned}$$

Gamma. Here we have

$$\eta = (-\beta, \alpha - 1)^T, \quad T = (x, \ln(x))^T, \quad h(x) = 1, \quad A = \log \Gamma(\alpha) - \alpha \log \beta$$

If $P = \text{gamma}(\alpha_2, \beta_2), Q = \text{gamma}(\alpha_1, \beta_1)$, where α_i, β_i are the so-called shape parameters, then

$$R(P \| Q) = \frac{\alpha_1}{\beta_1}(\beta_2 - \beta_1) + (\alpha_1 - \alpha_2)(\psi(\alpha_1)) - \log(\beta_1) + \log \frac{\Gamma(\alpha_2)\beta_1^{\alpha_1}}{\Gamma(\alpha_1)\beta_2^{\alpha_2}}.$$

Binomial. Here

$$\eta = (-\beta, \alpha - 1)^T, \quad T = (x, \ln(x))^T, \quad h(x) = 1, \quad A = \log \Gamma(\alpha) - \alpha \log \beta.$$

If $P = \text{binomial}(n, p_2), Q = \text{binomial}(n, p_1)$, then

$$R(P \| Q) = \log \frac{p_1^{\mu_1} (1 - p_2)^{\mu_1 - n}}{p_2^{\mu_1} (1 - p_1)^{\mu_1 - n}}, \quad \mu_1 = np_1.$$

Poisson. In this case

$$\eta = \log \lambda, \quad T = x, \quad h(x) = \frac{1}{x!}, \quad A = \lambda$$

If $P = \text{Poisson}(\lambda_2)$, $Q = \text{Poisson}(\lambda_1)$, then the relative entropy distance is

$$R(P \| Q) = \lambda_1 - \lambda_2 + \lambda_2 \log \frac{\lambda_2}{\lambda_1}.$$

Note that relative entropy formula is invariant under shifting and scaling of both distributions simultaneously. In fact, suppose that ψ and its inverse are both well defined and measurable, X and Y have distributions P and Q , and that $\psi(X)$ and $\psi(Y)$ have distributions \bar{P} and \bar{Q} . Then [2, Lemma E.2.1] $R(\bar{P} \| \bar{Q}) = R(P \| Q)$. A list of relative entropy formulas follows.

Distribution	$\mathbf{R}((\cdot)_2 \ (\cdot)_1)$
Gaussian	$\frac{1}{2\sigma_1^2} [(\mu_1^2 - \mu_2^2) + (\sigma_2 - \sigma_1)] + \log \frac{\sigma_1}{\sigma_2}$
Beta/Uniform	$(\alpha_2 - \alpha_1) [\psi(\alpha_2) - \psi(\alpha_2 + \beta_2)]$ $+ (\beta_2 - \beta_1) [\psi(\beta_2) - \psi(\alpha_2 + \beta_2)] + \log \frac{B(\alpha_1, \beta_1)}{B(\alpha_2, \beta_2)}$
Gamma	$\frac{\alpha_1}{\beta_1} (\beta_2 - \beta_1) + (\alpha_1 - \alpha_2) [\psi(\alpha_1) - \log \beta_1]$ $+ \log [\Gamma(\alpha_2) \beta_1^{\alpha_1} / \Gamma(\alpha_1) \beta_2^{\alpha_2}]$
Binomial	$\log [p_1^{\mu_1} (1 - p_2)^{\mu_1 - n} / p_2^{\mu_1} (1 - p_1)^{\mu_1 - n}]$
Poisson	$\lambda_1 - \lambda_2 + \lambda_2 \log \frac{\lambda_2}{\lambda_1}.$

References

- [1] R. K. Boel, M. R. James, and I. R. Petersen. Robustness and risk sensitive filtering. *IEEE Trans. on Auto. Control*, 3:451–461, 2002.
- [2] P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. John Wiley & Sons, New York, 1997.
- [3] P. Dupuis, M. R. James, and I. R. Petersen. Robust properties of risk-sensitive control. *Math. Control Signals Systems*, 13:318–332, 2000.
- [4] S.R.S. Varadhan. *Large Deviations and Applications*. CBMS-NSF Regional Conference Series in Mathematics. SIAM, Philadelphia, 1984.
- [5] D. Xiu. Efficient collocational approach for parametric uncertainty analysis. *Comm. in Computational Physics*, 2:293–309, 2007.

- [6] D. Xiu. Fast numerical methods for stochastic computations. *Comm. in Computational Physics*, 5:242–272, 2009.
- [7] D. Xiu and J. Hesthaven. High-order collocation methods for differential equations with random inputs. *SIAM Journal on Scientific Computing*, 27:1118–1139, 2005.
- [8] D. Xiu, J. Jakeman, and M. Eldred. Numerical approach for quantification of epistemic uncertainty. *Journal of Computational Physics*, 229:4648–4663, 2010.
- [9] D. Xiu and G. Karniadakis. The Weiner-Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24:619–644, 2002.